

PEG 2.0: Future-gazing through a socio-linguistic and historical lens^{*}

Verónica Dahl¹

Abstract

We discuss the importance of taking into account holistic historical and socio-linguistic perspectives to inspire us in clarifying what ends we should pursue- a discussion that is essential for the alignment that we are even asking the new forms of AI to have. We propose a few directions in that respect, which we feel are also conducive to redirecting those AI efforts that stress imitative or dominating goals, into the overall goal of *complementing* human intelligence to help people achieve shared prosperity sustainably through equitable, pro-social collaboration.

Keywords

Prolog Education, Logical thinking, Social Issues, Socio Linguistics, Holistic History, Verifiable Knowledge Bases

“Man is the most insane species. He worships an invisible God and destroys a visible Nature. Unaware that this Nature he’s destroying is this God he’s worshipping.” - Hubert Reeves, Canadian astrophysicist

1. Introduction

Since its inception, Prolog Education 2.0 was conceived as a world consortium in support of democratizing reliable thinking, coding and problem solving skills, for higher chances of overcoming the multiple crises that affect us all. Its aims are thus in great part societal, and its main means are the development, deployment and wide dissemination of Prolog-based teaching and learning tools conducive to our goals.

Yet the overall discipline where Prolog languages belong, namely AI, is not explicitly concerned with societal aims -although it is evident that its social effects are deep, numerous and wide-ranging, and although this fact has prompted some work along the lines of bringing the field to some "ethical", "responsible", "nurturing" or similarly worded point of fruition.

What mainstream AI has typically concerned itself with is how to pass the Turing test, i.e., how to *mimic the results* of human intelligence convincingly, by whatever means. This goal leaves the main contextual questions largely unaddressed: *AI for whom? by whom? for what? how?* Unstated, these questions are a tacit invitation for the status-quo objectives - or worse- to sneak in unanalysed and silently respond in its own interests.

With this article we hope to motivate discussion of such issues and their consequences to our field and to our PEG 2.0 endeavors. Section 2 provides historic context; Section 3 examines the role that societal systems in general play in determining technology’s goals; Section 4 discusses the role of contemporary social systems in determining the goals of Programming Languages and of AI and examines how the different contexts in which the two AIs developed have affected both their features and their goals; Section 5 discusses possible perceptions of logic as elitist in relationship to our field; Section 6 discusses the societal potential of AI; Section 7 summarizes how much of symbolic AI’s potential has already been exploited, with great success, by PEG 2.0; Section 8 provides my personal view of how we can do more, and Section 9 concludes by stressing our potential and need to redirect those AI efforts that

Joint Proceedings of the Workshops and Doctoral Consortium of the 41st International Conference on Logic Programming, September 9–13, 2025, Rende, Italy

✉ veronica_dahl@sfu.ca (V. Dahl)

🌐 <https://www.sfu.ca/people/veronica.html/> (V. Dahl)

🆔 0000-0002-1159-1374 (V. Dahl)



© 2025 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

stress imitative or dominating goals, into the overall goal of *complementing* human intelligence to help people achieve shared prosperity sustainably through equitable, pro-social collaboration.

While I do not claim that the views here expressed are representative of those of each and every PEG 2.0 member, I gratefully acknowledge the influence upon them that our assiduous meetings and discussions over the years exerted.

2. Historic Context

As thoroughly researched for instance by Riane Eisler [1], modern anthropology has allowed us to uncover interesting traits about ancient societies in what is now southern Europe: they were by and large based on mutual respect, caring, and cooperation; they represented divine power as benevolent Goddesses; and viewed power in general as the ability to create and support life. Not that violence or abuses, for instance, did not exist, but they were neither culturally accepted nor viewed as either societally or divinely ordained.

This eminently cooperative and gregarious nature- which studies of epigenetics attribute in fact to all our species as our innate characteristic ¹ - is what allowed us to develop, in order to share, technologies of production that still form the basis of all our modern technologies.

When the shift from partnership to dominator norms occurred, many thousand years ago, it split humanity into arbitrary hierarchies with women at the bottom, turned together with their children into mere male possessions. To achieve this extraordinary departure from human nature, it was key to de-throne the Goddess of Life that "primitive" peoples adored, and reconceptualize both Her and women in general as "less-than".

New values and a new, punitive God that upheld them were imposed through a millenia-lasting combination of indoctrination and violence. These new values were: forced, arbitrary hierarchies (vs. egalitarian social structures and institutions); male hegemony (vs. full human rights to all ²); a view of power as the ability to destroy (vs. viewing power as the ability to create and nurture); cultural acceptance of abuse and violence (vs. cultural rejection and intolerance thereof) and the naturalization of domination (vs. belief in mutual respect, caring and cooperation).

3. The role of social systems in determining technology's goals

The two above listed, opposing sets of values, which we could call *supremacist* vs. *equitist* ³, give rise to social systems respectively orienting to either domination or to equitable cooperation. The prevalence of one or the other has enormous implications for technology: societies where domination prevails emphasize *technologies of destruction* in order to dominate; whereas societies that orient towards equitable cooperation emphasize *technologies of production* in order to share equitably.

Obviously, the degree to which societies lean to each varies. Also, the above characterization is not to say that there is, for instance, no cooperation in domination-based societies, but it happens mostly in service of the bigger, domination-laden picture, e.g. a company's personnel might fully cooperate with each other, but do so in order to destroy the "competition" rather than to help achieve shared prosperity for all. Similarly, competition in equitable social systems may still exist, but will stress excellence and sharing rather than domination.

¹As researched e.g. in [2], departures from this collaborative nature are only achievable through the fear and force impositions needed for domination-orienting social systems to hold together

²Once the massive imbalance represented by male hegemony was successfully established globally after 2,000 years of relentless, violence backed indoctrination, further imbalances of power became more easily acceptable, growingly splitting also men into arbitrary hierarchies of domination. Hence our stress on "human rights **to all**".

³Terms such as "humanistic" having been co-opted to suggest "only men", we believe "equitist", less contaminated with double speak, does a better job of evoking partnership values

4. The role of contemporary social systems in determining AI's goals

Undeniably, societies worldwide still lean systematically towards the dominator values that were introduced so many millenia ago. Consider for instance that:

- *domination* (vs. belief in mutual respect, caring and cooperation) *has been naturalized* to the extent that inequitable rules of exchange have become the unquestioned "norm" everywhere
- societal endorsement of arbitrary hierarchies perpetuates e.g. male over-representation in decision-making positions while curtailing human rights for women, who globally enjoy only three quarters of the human rights allowed to men⁴ while being increasingly subjected to male violence⁵ and trapped into double workloads (one unpaid, one under-paid) misrepresented as "women's lib"⁶
- power is still mostly viewed as the ability to destroy, with economists counting death industries as "productive" while disregarding the essential contributions of unpaid carers
- cultural acceptance of abuse and violence is undeniable when considering world statistics of aggressions even during "peace".

AI languages being a type of programming languages, we shall start by examining the influence of social systems upon the PL field.

4.1. Social systems and the Programming Languages (PL) Field

In her recent talk and companion paper on "Programming for All: A Feminist Case for Language Design" ⁷, Felianne Hermans observed that, given that programming languages are interfaces between people and computers, the field could reasonably be expected to study its languages in relation to both, whereas in reality, research in dominant PL conferences focuses on coding, language features, quantitative methods and formal methods. This leaves the study of *users* of programming languages, of *applications* or of *analysis tools* in no person's land and prevents the PL community from having a more holistic view of programming language use.

4.2. PL Actors and their values

Hermans hypothesizes that PL's dominant culture prioritizes theory and formalism over people and social impact, limiting diversity. But the other way around is also the case: as she also observes, the range of acceptable knowledge creation tends to narrow to the implicit and explicit values of its actors.

Therefore, to understand these values, looking only at the societal values under which these fields were created is not enough. We must also look at the field's present demographics: both in PL in general and AI in particular, actors and deployers exhibit quite minimal demographic diversity- all the more shocking because at the field's inception, this was not the case.

Enter what has been described as the field's masculinization, aided in Hermans' view by Dijkstra's mathematization of it in order to, in his own words, "make the programming languages field more prestigious".

Indeed, the values that privilege difficulty per se, bragging rights, etc. over social value of whatever difficulty do tend to be exhibited most often by males in male-dominated societies, but we prefer to call them *supremacist* or *dominator* values, since gender alone cannot predict adherence: even those deriving personal benefit from an unequal status quo often rebel against it ⁸ in solidarity with others or

⁴<https://www.worldbank.org/en/news/press-release/2024/03/04/new-data-show-massive-wider-than-expected-global-gender-gap>

⁵<https://www.who.int/news/item/09-03-2021-devastatingly-pervasive-1-in-3-women-globally-experience-violence>

⁶While this misnomer subsists, female rights are receding worldwide (<https://www.downtoearth.org.in/governance/womens-rights-regressed-in-a-quarter-of-countries-in-2024-un-report-reveals>).

⁷https://drive.google.com/file/d/1JzXo6R3ol6EaNATTD-p8_GJF5Rrom-58/view?usp=sharing

⁸Conversely, dominator values can be embraced not only by dominators but also by dominated people, cf. phenomena like endo-racism and endo-sexism, where the dominated deny their predicament by trying to assimilate themselves to the ruling class through adopting the very values that keep them in fact subservient.

in recognition that their unearned privilege is not worth the (much less advertised) price they must pay, e.g. in terms of their own de-humanization or their providing free corpses for the expansionist adventures of some ruler. Further, not all men were or are particularly privileged: just more privileged than those even lower in the arbitrary hierarchies of the day. As modern times keep increasing the gap between those privileged and others and pushing the violence needed to dominate into suicidal levels, more men are realizing that even they have an interest to join those striving for equitism.

However, it is true that arbitrary hierarchies everywhere have placed domineering males at the top, so that their toxic values are often, in error, seen as typical of men in general. And it is these values that seem to have motivated the centuries-old devaluation of all things "soft" as feminine as "secondary" and unimportant. From the mathematization of the PL field onwards, Hermans argues, "soft" sciences were exited from the field and relegated to Human-Computer Interaction venues, less prestigious- and therefore, less influential- than PL publications. Other authors have amply documented the historic masculinization of not only the PL field but of Computing Sciences in general, e.g. [3].

4.3. The Two AIs and their different orientations

While our previous discussion applies to AI as well as to the LP field, sharp distinctions can be observed feature-wise between symbolic and subsymbolic AI.

Symbolic AI brings clarity and trust, being truth-explicit, logic based, explainable, verifiable, transparently executable, low resource and - most importantly- reliable: its results are trustworthy provided the (examinable) facts and rules it is fed also are. These features come **at a cost in terms of user-friendliness and amounts of knowledge available**.

Sub-symbolic AI brings pattern recognition and scalability, excelling as it does in uncovering patterns present in large amounts of data without having to understand them or verify their truth value. These features come **at a cost in terms of explainability, transparency, reliability, trustworthiness of results and ecologic footprint**.

4.4. The important, contextual questions

The distinctions above noted can help us answer the contextual question:

- **AI for what?** Symbolic AI is best suited for application-specific, data and inference-sensitive, truth-crucial or resource-crucial ends; whereas subsymbolic AI performs better at a) applications where there is an underlying theory whose rules can be captured with data, such as protein folding, or where enough good models can be consulted, such as in game playing or weather prediction and b) relatively blind but massive form-based exploration of a very specific kind of concepts: those deposited in the web- no matter if true or false, "neutral" or biased [4].

Other important questions on context are:

- **How achieved?** The lack of explicit societal planning re. this question is allowing resource allocation decisions to be made in disregard of obvious dire consequences to numerous people- e.g. in terms of drinking water being diverted into subsymbolic AI's consumption needs. While symbolic AI can also be put to inequitable uses, it does not depend as Big AI does on as vast an exploitation of resources like water, energy or minerals, nor on the vast amounts of the human unpaid work daily contributed -often unwittingly- by users worldwide.
- **AI for whom?** Symbolic AI is less accessible than Subsymbolic AI, into which formidable amounts of research and money have poured. Undoubtedly all of us that use from a cell phone's completion verbots, poor as they may be, to most helpful daily consultation of chatbots, benefit in one way or another from Big AI's ubiquitousness. The financial benefit of Big AI, however, goes mostly to the handful of companies that have taken control of our personal and other kinds of data, while largely depending on unpaid invisible labor [5, 6]. Generalized accessibility is one of the achievements that PEG 2.0 is working on for Symbolic AI.

- **AI by whom?** Currently, subsymbolic AI's main deployers are a small group of companies and executives intent upon automating human tasks to their advantage, in ways that often breed misinformation and manipulation, sideline humans, exacerbate inequality, are unsustainable and enable those who control AI technologies to rule over the rest of us and exploit our collective work, applying it in particular to "replace" human workers (a misnomer, since AI's shortcomings with respect to human workers laid off must often be compensated by the public's unpaid time and work, and only human workers contribute to tax, social security or benefits). Symbolic AI's actors, while also being demographically unrepresentative, are largely world scientists and research institutions preoccupied with scientific soundness- and who often have societal good explicitly in mind. However, they do not have easy access to the formidable amounts of money and resources of all types needed to even train small BigTech models. Consequently, the number of Machine Learning publications produced is increasingly being funded or authored by corporate affiliated representatives rather than by academics [7].

4.5. Does "AI Ethics" help promote equity or domination?- a language-based analysis

This question is a useful one to explicitly ask of any discipline, social system or proposed measure, since answering it can help clarify what is essential and often hidden behind "-isms" or similarly loaded labels that abound in our era of confusing double speak.

We include in "AI Ethics" (also called "Good AI", "Responsible AI". etc.) all the literature devoted to defining the values and ideas that should guide AI advances and deployment. This body of literature boomed in 2017-19, when transparency became important after impressive societal harms made possible by AI (such as electoral manipulations by Cambridge Analytica, or the first human killing by Uber self-driving cars) shocked public opinion.

Surprisingly, the subject matter is never defined in many of these works, so researchers have been forced to analyse its salient language in order to characterize it.

Jobin et al [8] examined 84 documents containing ethical guidelines for intelligent autonomous systems, some of which were at the time the most cited guidelines in the literature. It detected as the most common principles: transparency (86% recurrence), **justice and equity (81 %)**, non-maleficence (71 %), responsibility (71 %), privacy (56 %), accountability, beneficence (48 %), freedom and autonomy (40 %), trustworthiness (33 %), **sustainability (16 %)**, **dignity (15 %)** and **solidarity (7 %)**.

Of these ethical principles cited, the five most recurrent ones, namely transparency, justice, non-maleficence, responsibility, and privacy have been corroborated as such in two more studies: [9], which selected 21 documents deemed relevant in the international discourse (IEEE, Google, Microsoft, and IBM), and [10], which studied 200 governance policies and ethical guidelines for AI usage published by public bodies, academic institutions, private companies, and civil society organizations worldwide.

[10] also found that only 55.5 % of documents seek to define what is the object of their discourse- a fact made even more worrisome given that there is no consensual definition of what "artificial intelligence" is and what it is not-; that most of the documents only prescribe normative claims without spelling out the means to achieve them; and that **the overwhelming majority of government documents (91.6 %) opt for "soft" forms of regulation, with no means to enforce them**. As a result, the private sector centers their interest in Ethical AI around the application of "technical" solutions to social problems, or simply around evading regulation.

These flaws alone suggest that AI Ethics is too brittle a field to be able to effectively rescue AI from the domination contexts it might be embedded in. In addition, the studies mentioned show that *the principles that could best propel us fully back into equitable societies of cooperation (namely **solidarity and sustainability**) are hardly even mentioned in the literature, with the latter having grown in mentions in the past few years, but still shockingly under-addressed*.

An interesting exception re. solidarity, however, is emerging: studies from Latin America, like [11], show that most of the most cited values are equity-promoting, although sustainability is also neglected: **fairness/non-discrimination (100 %)**, privacy (97 %), accountability (97 %), transparency/explainability (94 %), safety/security (81 %), professional responsibility (78 %), human control of technology (69 %),

and **promotion of human values (69 %)**.

4.6. Does "AI Ethics" help promote equity or domination?- a content-and-intent-based analysis

Relational ethics analyses of machine-learning based AI such as [12, 13],⁹ yield the following important points:

Content-wise,

- Big AI performs best on phenomena of low ethical stake. The more socially complex a problem is, the less capable machine learning systems are of capturing the phenomena in a model, and the greater the likelihood becomes of them causing societal damage.
- While "AI for good" is being pushed even by large organizations such as the UN or by humanitarian places intent on using AI for political or social ends, most such initiatives tend to have little value for the actual communities they are supposed to serve.
- A lot of subsymbolic AI is not about understanding the historical data, but about using it to predict the future, hence it builds models that are best at replicating historical patterns. In consequence, underlying social and political issues undergo crippling simplifications and the models often end up potentiating biases, discrimination, harm.

Intent-wise,

- The social infrastructure of Big AI itself contradicts social good, since the entire ML field is based on the business model of maximizing profit as its main goal, with scant concern for social welfare.

Birhane concludes that even in "AI for social good", *Big AI tools are unlikely to be helpful*, and calls for a more accountable, human-centric, and context-aware approach to "social good".

5. Is logic elitist?

Our analysis in Section 4.3 of features of the two types of AI suggests that our kind of logic, namely executable, is largely exempt from the elitist intent that Hermans traces historically to the introduction of logic into the PL field: that of making the field more difficult- and therefore, more prestigious- partly *by requiring a focus on logical proofs* that detracted from focussing on social aspects.

Our field does not use logic for mandating proofs for prestige-buildup. Not that proofs are unimportant to us, but they are mostly automated and hidden, the intent being to make things *simpler*, rather than more difficult, for the user. Even the form of logic we use - Horn clause logic- is simplified. But does our aim of democratizing access to logical thinking constitute enough of a focus on social aspects? The demographic composition of our actors is no more inclusive than that of PL in general, which suggests we might be missing valuable social perspectives.

The next interesting question, therefore, is: can we actually *help recover the focus on social aspects* that the PL field lost? We next examine this question, in the trust that most readers will agree that such a refocussing is desirable and even needed.

6. AI's Societal Potential

6.1. Symbolic AI's Societal Potential

Arguably, the language features our field is fairly unique in providing (truth-centered, reliably executable, examinable, provable, traceable, trustworthy, transparent, accountable, verifiable, logical) are the same features that societies orienting towards equity must embrace in order to achieve their goal.

⁹<https://www.youtube.com/watch?v=KbUmEtvS0vE>, <https://aiforgood.itu.int/event/ai-for-social-good-the-new-face-of-technosolutionism/>

Of these features, **truth-centerdeness** is probably the most important one, since reasoning about wrong data cannot give good results no matter how good the reasoning skills applied are. We are already embracing it.

Another necessary feature, presently either absent or inexplicit, but necessary if we are to clearly adopt a focus on social aspects, is **the overall goal of developing and placing our tools at the service of societal equity and shared prosperity**, because not even total access to truth and to superb reasoning skills can bring about better societies if placed at the service of hierarchies of domination.

As the next section will show, PEG 2.0 already largely aligns with this objective.

6.2. Sub-symbolic AI's societal potential

Big Data AI's features, in particular its **excellence in examining and uncovering patterns in large amounts of data, have already enabled major breakthroughs** in areas like healthcare, accessibility, or climate science, in which it can be complemented by good, already existing models. It can even be helpful in more general tasks, e.g. as a research assistant (to develop summaries, write code, explain concepts, etc.)- as long as its departures from truth can be caught and fixed case by case by the user. **The key is to endow users with a well-rounded education** that enables them to ask the right questions and interpret the results intelligently.

The present lack of universal availability of such skills and of education in general, coupled with Big AI's incursions into applications that are either clearly supremacist or for which it is clearly not ripe, has sometimes yielded catastrophic results such as biased recommendations or hate speech amplification- particularly when data repositories of dubious quality were unduly relied upon.

Therefore, it is even more important for Subsymbolic AI to explicitly adopt the feature we mentioned as necessary to Symbolic AI as well, namely **the overall goal of developing and placing our tools at the service of societal equity and shared prosperity**.

If equity had been BigAI's overall goal from inception, we would not have used it for societal applications where the consequences of potential error can be dire, as in face recognition, nor for those where the intended applications have devastating societal consequences for some group, such as Denuding AIs.

We would have instead placed Big AI's excellence at the service of e.g. anthropologists, sociologists and planners, who could have used it to **detect and study** bias, hate speech or political manipulation by either real or counterfeit social media users, **in the aim of overcoming such biases and manipulations** rather than amplifying them. Authors like Paolo Rosso and Roberta Calegari have devoted many years to this kind of research ¹⁰.

7. PEG 2.0's Societal Impact

Our consortium of just a few dozen of volunteer experts has paved the way on many fronts, in only three years and despite scant resources, to spreading logical reasoning and problem solving skills worldwide.

The initiatives that led to this collective honour can be grouped into the following categories:

- **Integrating Logic Programming with other STEM or non-STEM disciplines, addressing students and teachers:** Two projects led by *Yuanlin Zhang* (*Logic Programming for K-12 students* and *Logic for Data Science*) and their corresponding theoretical frameworks have been developed and tested for using Logic to integrate the data science foundations across topics and grades in math, computing and statistics. Integrations with non-STEM disciplines span all levels and center mostly around artistic design [14], language and societal planning [15, 16, 17, 18]. Interestingly, these initiatives privilege *generative* applications of LP.

¹⁰<https://scholar.google.com/citations?user=HFKXPH8AAAAJ&hl=en>, publications

<https://www.unibo.it/sitoweb/roberta.calegari/>

- **Introducing Logic Programming into elementary schools:** Besides the elementary school component of the above noted initiative, under the leadership of *Laura Cecchi*, courses for 9 and 10 year olds were developed and tested in ten successful experiences in Neuquén, Argentina. This initiative had the support and participation of the Dirección de Educación Digital of the Neuquén Province, and yielded ten experiences at a total cost of 250 dollars. Courses for elementary school teachers were developed and tested in conjunction, aided by student volunteers specifically trained into sensitivity to racial and gender issues.
- **Creating interdisciplinary knowledge bases:** The initiative "Digital Bulgaria in Prolog", led by Veneta Tabakova-Komsalova, Magda Maglijanova and Asya Stoyanova-Doycheva, initially aimed at introducing Prolog logic programming to students at grades 5-7 (junior high school) and grades 8-12 (senior high school)- was extended to even younger students. It involves a network of schools that incorporate logic programming into their curricula, and have initiated pilot programs in Burgas and Plovdiv, leveraging interdisciplinary examples from Bulgaria's cultural and historical heritage and integrating subjects such as biology, chemistry, history, and geography;
- **Zoom-accessible courses, materials and activities for teenagers and young adults.** These include:
 1. **Reasoning Hackathons**, led by *Gopal Gupta* at UT Dallas and in coordination with various European universities, have capitalized on the popularity among the young of Software hackathons by introducing Logic Programming in the content to compete about, thereby transforming them into reasoning competitions;
 2. **Logical Language**, led by *Bob Kowalski* at Imperial College and *Jacinto Davila* in Venezuela, provides facilities for easier expression of logic programs using natural languages such as English, Spanish and French; *Logic for all* pioneered by *Michael Genesereth* at Stanford, provides a novel approach to teaching logic which is offered online for free in perpetuity, plus summer camps for high school students and International Logic Olympiads for secondary school students. Michael is at present further specializing this approach specifically into LP and Prolog;
 3. **Exploring Generative Logic**, EGL, an educational program specifically tailored for students in art and design, was developed by *Christian Jendreiko* at the University of applied sciences, Duesseldorf, demonstrating a new didactic method of learning logical thinking through creative action, which has great potential to provide diverse target groups with an accessible entry point to logical thinking. The diversity-integrating potential of this work is enormous and potentially game-changing;
 4. **Materials using active logic documents** are being developed at IMDEA Software and University of Evora respectively, by *José Morales* and *Salvador Abreu*, aimed at adapting existing Prolog courses to fully browser-side interactivity;

Rigorous research plus practical work too vast to explain here underlies, of course, all of these initiatives and many more not yet incorporated into our repertoire of web-accessible materials. Noteworthy among these are efforts by *Gopal Gupta* and *Paul Tarau* for integrating Big Data AI into societally useful teaching efforts and/or complementing its form-oriented ways with some kind of logic guidance (e.g. [19, 20]).

8. Can we do more?

Yes! My view of how includes: a) through paying special attention, in our activities and projects, to how uses of language and exposure to supremacist stereotypes will typically boycott socially positive endeavours; b) through redirecting those AI efforts that stress imitative or dominating goals, into the overall goal of complementing human intelligence to help people achieve shared prosperity equitably, within planetary limits. From our expertise's point of view, this involves centering truth to counteract the deluge of misinformation, fakes, impersonations, etc. we receive daily.

8.1. The role of language in promoting pro-social values

Much of the persistence of dominator values owes to how human languages are twisted into disseminating those values in the most effective way possible: by hammering them surreptitiously into our collective unconscious, day-in, day-out. This has proceeded mostly through language [21], through institutional pressure such as gender-biased judicial systems [22] and through supremacist stereotyping [23].

For instance L'Académie Française, in charge of defining the French language, forbade the feminine form of words denoting prestigious professions (such as *peintresse*, *doctoresse*, *magistrate*, *écrivaine*, *présidente*, *sculptresse*), allowing the feminine form only for subservient occupations such as "servante". This indoctrinated public opinion against women's rights, preparing the ground for those professions being subsequently made illegal for women. While this injustice was (much) later reversed, its linguistic roots remain, subtly but relentlessly repeating the message that women "do not quite belong" in the human category- a message also enthusiastically propagated by popular insults, e.g. those that redirect anger towards the rightful recipient's *mother*. The ground keeps perpetually being prepared, through such ubiquitous aggressions only unconsciously perceived, for the next wave of legal deprivations a dominator society might be preparing.

But just as language can be twisted into dominator values through coercion leading to subsequent adoption of those values, we can revert them into equitist values through most of us being conscious of, and proactive about, our own production and decodings of language.

As educators, we can for instance use disciplines such as semiotics to help students detect and deactivate toxic or double-speak language. PEG 2.0 could develop and provide to students and teachers handy tools for doing so, e.g. knowledge bases about the sememes and implicit meanings contained in important words, expressions, discourses or dialects and jargons such as teenage jargon, legalese, political speeches, etc.

As researchers, we can engage in hybrid, symbolic AI-led applications that actually put subsymbolic AI's excellence in finding patterns into the service of detecting and calling out discriminating patterns and even helping disseminate constructive ones instead. For instance, negative adjectives with common sexist connotation (e.g. "a nag") can be flagged out for non-sexist correction (e.g. into "opinionated").

8.1.1. The role of thoughtful AI jargon

A particularly important linguistic task for us AI-ers is to de-personify AI. Companies will refer to AI itself as "intelligent", "good", "responsible", "ethical", "nurturing", etc., as though it were a person, either for advertising ends, or to de-responsibilize themselves from some of AI's harmful effects, or to conceal how much human intelligence AI is using behind the scenes in order to hide its shortcomings where "intelligence" is concerned [24].

Just as personhood granted legally (while inaccurately) to companies resulted in humans being prevented from suing them for damages really done by them, speaking of AIs as though they were people paves the way to eventually granting them "human" rights, which would leave real humans even more resourceless than now when machines trample on their rights [25].

Let's resist the linguistic reflex -we all have it, since it's been instilled in us- **of referring to AI as if it were human, or even intelligent in the human sense**, which is embodied, feeling-informed, introspective, explainable, capable of reasoning, understanding-grounded, cooperative, relational and caring.

8.2. The role of de-stereotyping

Once instated, supremacist values needed not only loaded language impositions plus the inordinate amounts of violence it took for them to constantly survive and replicate (e.g. 400 centuries of witch burnings) but also, most importantly, forced and fallacious stereotyping: the natural humane values that stood in the way of domination had to be demoted, over and over as they kept reappearing century after century. What better way than to misrepresent them as "negative", "weak", "unimportant", and

"sub-human", therefore "inferior", and ALSO stereotype them as feminine, thus sealing the coffin around the fate of women and that of the Goddess that had symbolized Life, Love, and Nurturing? In complement, dominator values were misrepresented as typically male, "positive", "powerful", "important" and "fully human", therefore "superior", thus sealing around the fate of men an invisible coffin of abuse-backed cultural suppression of feeling and humaneness, the better to manipulate them into war and other supremacist ends.

Belief in these mutually reinforcing stereotypes actually harms men as well as others, although in different ways [26], but nevertheless, unfortunately, ensures their -active or by inertia- collaboration with the subordination and pauperization of others that is still essential to male domination.

However, it is the acceptance- particularly the unconscious acceptance- of stereotypes that determines which arbitrary hierarchies, new or old, are accepted. Therefore, **just as with language, it is important to develop educational tools that can help people see through stereotypes and distinguish natural hierarchies** (e.g. the right of medical doctor, above that of everybody else, to treat disease) **from arbitrary hierarchies of domination** (e.g. the "right" of a white man to higher wages, for work of equivalent value, than other people, such as of colour).

8.3. The role of humaneness and of consent: our professional responsibility

We put forward that **as AI experts, we have both the right and the responsibility** to a) refuse to participate in creating or helping deploy obviously criminal AIs (such as AIs designed to scout the ground exhaustively killing anything that moves, e.g. children or adult civilians), and b) demand that consent from those affected be sought before we participate in technical developments or deployments with obvious harmful side effects to groups or communities.

The latter duty is especially important given the already mentioned lack of proportionate representation that characterizes our field's actors. It is doubtful, for instance, that denuding AIs, which target mostly females and have obvious and unfortunately already verified victimizing potential, would have seen the light if sufficient girls and women had been consulted- or even proportionally represented among the designers and their hirers. As another example, had the peoples suffering drinking water shortages been consulted about how much of this water should be allocated to Big AI instead of to humans, chances are we would be redoubling efforts towards low-resource AIs ¹¹.

8.4. The leading role of Symbolic AI

As mentioned, collaborations between the two AIs are already under way, in the intent to harness the combined power of automatically discovering patterns in vast amounts of form and of automatically processing meaning representations and reasoning over more focussed information. Since the first is not sustainable, it stands to reason that its mega-resources approach will give way to more contained special applications whenever possible, leaving vast searches for cases in which it is truly needed. Already the addition of models to otherwise blind Big AI's search is complementing it fruitfully. Symbolic AI's aptness at low-resource reasoning power needs, in our view, to take the lead. We put forward that it is time to center those efforts trying to combine the two types of AI around reason, leaving Big AI's searches to be called as helpers when needed, but not leaving them to run the show, particularly when the social stakes are high or when historical patterns or current biases are precisely what we do NOT want to replicate.

8.5. The complementary role of Subsymbolic AI

Big AI, if driven by logical AI and implemented with equity and sustainability in mind, can be used to help solve socio-ecological problems such as lacks in human rights as democracy and equity, or in planetary rights such as CO2 overshoot.

¹¹ According to the United Nations Environmental Report, almost half of the world's population will by 2030 be facing severe water stress.

For instance, instead of being used to replace workers, it can be regulated into a) reducing their work-day thanks to the efficiency gained, thus freeing their time for the social service that will increasingly be needed, or giving them the opportunity for lifelong education or other societally useful activities; b) paying taxes, benefits, etc. just like human workers do; c) where it does make sense to lay workers off, helping them find and retrain into the green or caring jobs for which they will still be needed. It could also be used to detect hate speech, incitations to violence, unfair wages, etc. in the aim of eradicating them altogether.

As another example, Big AI can be used to remove the practical barriers that impede us from exercising democracy directly: a (non-capturable) social media platform could help us move into true democracy by reliably counting votes on thoroughly informed important issues, e.g. what proportion of our budget should go to the war industry, to universally free education, to public health and services? Public protest would not need demonstrations but simple vote counting. But of course, guarantees to respect and follow up on the winning motions should be in place, and AI's ecosystem should be cleaned up first: no fake accounts, no fake news, thoroughly verifiable, quality information, etc.

8.6. The role of truth: Worldwide repositories of verifiable knowledge

Perhaps the most societally influential action we can take is to counteract the flood of disinformation and fakes we are daily bombarded with, through developing, one knowledge base at a time, a commons-owned, free and not capturable worldwide knowledge network of documentedly verifiable truths.

The Bulgarian experience we have described in Section 7 exemplifies. Students go collect the data on whatever part of the knowledge base their team is involved in creating, and this information is checked for veracity by the university professors and specialists in charge of the experience, before being entered. Our extension of this idea would involve documenting the sources of every piece of information on a given theme (e.g. "the following 20 scientific articles, on pages such and such", or "newspaper X", with details including the journalistic standards enforced on its articles¹²).

This verified-truth approach, if also diversity-conscious, would by the way help us catch on time any dominator-induced gaps in data collection, e.g. asking who the drug Thalidomide had been tested on would have yielded, just as it did in the U.S. where the question WAS asked ¹³, the information (namely: only on men) that would have prevented also in the rest of the world the birth of children with no legs or arms. This approach would throw light upon the skewed priorities of dominator societies, at the same time as making the necessary questions evident to ask.

9. Conclusion

If PL/AI's dominant culture prioritizes theory and formalisms over people and social impact, it is because that culture is embedded in (plus demographically most representative of) our dominant societal culture, whose eternal quest for the ever-growing "power" of the few prioritizes for instance economic and politics over people and over science- a situation which has already led us to transgress most safe planetary limits.

As logical reasoning experts intent upon educating the entire world, as detentors of the sole type of AI which is trustworthy and can actually reason verifiably, **we owe it to the world to start focusing on how our tools can best be used to help communities propitiate the cultural shifts needed while simultaneously developing the resilience that they urgently must develop.** As we have argued, these shifts crucially involve **joining efforts around two main values: solidarity** (so that basic human needs- as per United Nation's Declaration of Human Rights ¹⁴- are taken care of for all

¹²This is particularly important in an era in which journalists are no longer required, as they were before, to prove that they had thoroughly checked the veracity of their assertions before these got published

¹³Anecdotically, it was asked by a woman, Frances Kelsey, who attained the position of reviewer for the U.S. Food and Drug Administration thanks to the felicitous confusion between "Francis" and "Frances" of the professors who admitted her into graduate studies in the belief that she was a man

¹⁴<https://www.un.org/en/about-us/universal-declaration-of-human-rights>

humans), **and sustainability** (so that planetary limits -as defined by Earth scientists [27]- are not further transgressed), for a chance of our race's survival.

Many fields, e.g. Economics [28], are contributing already to such ends. If from our field we can also help us all evolve into the universal adoption and manifestation of solidarity and sustainability as our guiding goals, we will be helping redirect AI to ends less "prestigious" than persuasively posing as intelligent, into the pressing, prestigious or not, goals of ***complementing human intelligence to help us all achieve shared prosperity through equitable cooperation***, supporting in particular the new tasks and skills needed to surmount our life-threatening socio-ecological crises.

As we hope to have shown, cultural shifts crucially involve language shifts, since it is mostly through language that we communicate our (conscious or unconscious) values. The cultural shift we need also involves eradicating manipulative ones necessary to perpetuate domination mindsets. AI tools can be developed to help people restore truth as a value, detect the toxic messages by which we die that language keeps encapsulating, and rationally analyse the roles that supremacist stereotypes play in our cultures.

Last but not least: we must engage in a wide-ranging campaign of accessible dissemination, extend our activity more purposefully beyond the academic world and into the "real" world, addressing the masses through Op-Eds, magazines, books, talks to the layperson, etc. We are already halfway there through our tools, resources and initiatives -in particular, that of teaching teachers worldwide- that are already changing the game in several countries. We need to reach many more.

Maybe not everything is lost. New generations are entitled to the help of the older generations, which have created the serious problems they alone will be left to solve. Let's imagine our way into as much of such a help as possible.

Dedication This article is warmly dedicated to the memory of my brother, the late Henry Dahl, whose lifelong quest for social justice greatly influenced me; to Yolanda Mansilla who fuelled my interest in math in elementary school; to the Faculty of Philosophy and Letters of Universidad de Buenos Aires and to my former professors there: the late Gabriel Bès and Alfredo Hurtado, who taught me Chomskyan Syntax and later became my collaborators; Beatriz Lavandera, from whom I learned Sociolinguistics; Gregorio Prieto who taught me Semiotics; Celia Jakubowicz, my professor of Psycholinguistics; and to my parents's memory: Selva Otero, for modeling both courage when powerless and equitism when in positions of power; and Ivar Dahl, whose informal teachings, including on French, Spanish and English Phonetics highly enriched my cultural sensitivity, fuelling my desire to build bridges between humanistic and formal sciences. Last but not least, to the memory of my thesis supervisor Alain Colmerauer, whose brilliance continues to inspire me.

Acknowledgments

My unending, big and warm THANKS goes to all participants past and present of PEG 2.0, for their inspiring discussions and generous devotion of many years constantly working for our dreams to come true. I am also most grateful to Laura Cecchi, Jacinto Dávila-Quintero, Bharat Jayaraman, Christian Jendreiko, Bob Kowalski, Veneta Tabakova-Komsalova, and Alvaro Videla for fruitful discussions on many of this article's themes, and to Evgeny Skvortsov, Raúl Carnota and the anonymous reviewers for their useful feedback on this paper's first draft. I am indebted to Argentina and to France for having crucially covered the cost of, respectively, my under-graduate and my graduate studies, on the sole condition that I excel studying, and to NSERC, for its support from grant 31611021.

Declaration on Generative AI

During the preparation of this work, the author used Gemini for grammar and spelling checks. After using this tool, the author reviewed and edited the content as needed, and takes full responsibility for the publication's content.

References

- [1] R. Eisler, *The Chalice and the Blade: Our History, Our Future—Updated With a New Epilogue*, HarperOne, 2011.
- [2] R. Eisler, D. P. Fry, *Nurturing Our Humanity: How Domination and Partnership Shape Our Brains, Lives, and Future*, Oxford University Press, Oxford, 2019.
- [3] E. V. Oost, Making the computer masculine, in: *IFIP TC9/WG9.1 Seventh International Conference on Woman, Work and Computerization: Charting a Course to the Future*, 2000, pp. 9 – 16.
- [4] A. Birhane, V. Prabhu, E. Kahembwe, Multimodal datasets: misogyny, pornography, and malignant stereotypes (2021). doi:10.48550/arXiv.2110.01963.
- [5] M. Gray, S. Siddharth, *Ghost Work: How to Stop Silicon Valley from Building a New Global Underclass*, Houghton Mifflin Harcourt, 2019.
- [6] C. D'Ignazio, L. F. Klein, *Data Feminism, Strong Ideas*, MIT Press, Cambridge, MA, 2020.
- [7] A. Birhane, P. Kalluri, D. Card, W. Agnew, R. Dotan, M. Bao, The values encoded in machine learning research, 2022. URL: <https://arxiv.org/abs/2106.15590>. arXiv:2106.15590.
- [8] A. Jobin, M. Ienca, E. Vayena, The global landscape of AI ethics guidelines, *Nature Machine Intelligence* 1 (2019) 389–399. URL: <http://dx.doi.org/10.1038/s42256-019-0088-2>. doi:10.1038/s42256-019-0088-2.
- [9] T. Hagendorff, The ethics of AI ethics: An evaluation of guidelines, *Minds & Machines* 30 (2020) 99–120. URL: <https://doi.org/10.1007/s11023-020-09517-8>.
- [10] N. K. Corrêa, C. Galvão, J. W. Santos, C. Del Pino, E. P. Pinto, C. Barbosa, D. Massmann, R. Mambrini, L. Galvão, E. Terem, N. de Oliveira, Worldwide ai ethics: A review of 200 guidelines and recommendations for ai governance, *Patterns* 4 (2023) 100857. URL: <https://www.sciencedirect.com/science/article/pii/S2666389923002416>. doi:<https://doi.org/10.1016/j.patter.2023.100857>.
- [11] N. Fjeld, J. and Achten, H. Hilligoss, A. Nagy, M. Srikumar, Principled artificial intelligence: Mapping consensus in ethical and rights-based approaches to principles for ai, Berkman Klein Center for Internet & Society, 2020.
- [12] A. Birhane, Algorithmic injustice: a relational ethics approach, *Patterns* 2 (2021) 100205.
- [13] A. Birhane, P. Kalluri, D. Card, W. Agnew, R. Dotan, M. Bao, The Values Encoded in Machine Learning Research, in: *ACM Conference on Fairness, Accountability, and Transparency*, ACM, New York, NY, USA, 2022.
- [14] C. Jendreiko, Generative logic: Teaching prolog as generative ai in art and design, in: *PEG 2.0 Workshop*, 2024.
- [15] V. Dahl, Doughnut computing: Aiming at human and ecological well-being, in: *6th Int. Conference on the History and Philosophy of Computing (HAPOC-6)*, 2021.
- [16] V. Dahl, J. J. M. Navarro, Doughnut computing in city planning for achieving human and planetary rights, in: *IWINAC 2022: 9th International Work-Conference on the Interplay Between Natural and Artificial Computation*, IWINAC, Tenerife, Canarias, 2022.
- [17] V. Dahl, G. Bel-Enguix, V. Tirado, E. Miralles, Grammar induction for under-resourced languages: The case of ch'ol, in: J. Gallagher, R. Giacobazzi, P. Lopez-Garcia (Eds.), *Analysis, Verification and Transformation for Declarative Programming and Intelligent Systems*, number 1316 in LNCS, Springer, 2023.
- [18] V. Dahl, L. A. Cecchi, Introducing Prolog in language-informed ways, in: D. S. Warren, V. Dahl,

- T. Eiter, M. Hermenegildo, R. Kowalski, F. Rossi (Eds.), Prolog - The Next 50 Years, number 13900 in LNCS, Springer, 2023.
- [19] P. Tarau, Leveraging LLM reasoning with dual horn programs, in: E. Erdem, G. Vidal (Eds.), Practical Aspects of Declarative Languages - 27th International Symposium, PADL 2025, Denver, CO, USA, January 20-21, 2025, Proceedings, volume 15537 of *Lecture Notes in Computer Science*, Springer, 2025, pp. 163–178. URL: https://doi.org/10.1007/978-3-031-84924-4_11. doi:10.1007/978-3-031-84924-4_11.
 - [20] G. Gupta, H. Wang, K. Basu, F. Shakerin, E. Salazarand, S. C. Varanasi, P. Padalkar, S. Dasgupta, Logic-based Explainable and Incremental Machine Learning, in: D. S. Warren, V. Dahl, T. Eiter, M. Hermenegildo, R. Kowalski, F. Rossi (Eds.), Prolog - The Next 50 Years, number 13900 in LNCS, Springer, 2023.
 - [21] R. Lakoff, The Language War, University of California Press, California, 2020.
 - [22] The Gender of Constitutional Jurisprudence, Cambridge University Press, 2004.
 - [23] P. Bourdieu, Language and Symbolic Power, Polity Press, Maiden, MA 02148, USA, 2020.
 - [24] M. Gray, S. Suri, Ghost Work: How to Stop Silicon Valley from Building a New Global Underclass, Houghton Mifflin Harcourt, 2019.
 - [25] A. Birhane, J. van Dijk, Robot Rights?: Let’s Talk about Human Welfare Instead, in: Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society, ACM, 2020, p. 207–213.
 - [26] L. Plank, For the Love of Men: From Toxic to a More Mindful Masculinity, Bridgewater State College, 2021.
 - [27] W. Steffen, Å. Persson, L. Deutsch, J. Zalasiewicz, M. Williams, K. Richardson, C. Crumley, P. Crutzen, C. Folke, L. Gordon, et al., The anthropocene: From global change to planetary stewardship, *Ambio* 40 (2011) 739–761.
 - [28] K. Raworth, Doughnut Economics: Seven ways to think like a 21st-Century Economist, Chelsea Green, White River Junction, Vermont, 2017.